

On the physical interpretation of a meta-analysis in the presence of heterogeneity and bias: from clinical trials to Mendelian randomization

Jack Bowden^{1,2*}

¹ *MRC Integrative Epidemiology Unit, University of Bristol, UK.*

² *MRC Biostatistics Unit, Cambridge.*

Chris Jackson

MRC Biostatistics Unit, Cambridge, UK

Summary. The funnel plot is a graphical visualisation of summary data estimates from a meta-analysis, and is a useful tool for detecting departures from the standard modelling assumptions. Although perhaps not widely appreciated, a simple extension of the funnel plot can help to facilitate an intuitive interpretation of the mathematics underlying a meta-analysis at a more fundamental level, by equating it to determining the centre of mass of a physical system. We used this analogy, with some success, to explain the concepts of weighing evidence and of biased evidence to a young audience at the Cambridge Science Festival, without recourse to precise definitions or statistical formulae. In this paper we aim to formalise this analogy at a more technical level using the estimating equation framework: firstly, to help elucidate some of the basic statistical models employed in a meta-analysis and secondly, to forge new connections between bias adjustment in the evidence synthesis and causal inference literatures.

Keywords: Meta-analysis, Funnel plot, Causal inference, Estimating equations, Egger regression, Mendelian randomization.

1. Introduction

Let y_i , $i = 1, \dots, k$, represent summary estimates of the same apparent quantity from k independent information sources in a meta-analysis. The i 'th estimate is associated with a fixed and known variance, s_i^2 . The standard fixed effect model assumes

$$y_i = \mu + s_i \epsilon_i, \quad \epsilon_i \sim N(0, 1) \quad (1)$$

The focus for inference, in terms of point estimates and confidence intervals is the population mean effect parameter μ . The fixed effect estimate $\hat{\mu}$ and its variance are given by the well known formula

$$\hat{\mu} = \frac{\sum_{i=1}^k w_i y_i}{\sum_{i=1}^k w_i}, \quad \text{Var}(\hat{\mu}) = 1 / \sum_{i=1}^k w_i. \quad (2)$$

Meta-analysis has a long history in medical research, where the information source has traditionally been a randomized clinical trial, comparing an experimental treatment against standard therapy. However, its use spans the entire scientific spectrum, for example in the areas of psychology (Hedges, 1992); inter-laboratory experiments (Paule and Mandel, 1982); particle physics (Baker and Jackson, 2013); Mendelian randomization (Burgess and Thompson, 2010) and many others.

Regardless of the application area, if s_i and ϵ_i are mutually independent, then equation (2) will give a consistent estimate for μ . In order to provide a convenient shorthand that easily generalizes to more complex models, we write $\perp\!\!\!\perp_{(s_i, \epsilon_i)}$ to denote this assumption.

1.1. *The funnel plot and the fixed effect model*

The funnel plot (Sterne and Egger, 2001) is a graphical visualisation of data from a meta-analysis. It plots study i 's effect size on the x-axis versus a measure of precision on the y-axis (usually $1/s_i$). The independence assumption $\perp\!\!\!\perp_{(s_i, \epsilon_i)}$ of model (1) implies that there should be no correlation between effect size and precision. If this is the case, then the plot should appear symmetrical, in that the results from smaller studies should funnel in towards those from larger studies. It therefore provides a simple means to visually assess the plausibility of $\perp\!\!\!\perp_{(s_i, \epsilon_i)}$.

Although perhaps not as widely appreciated, a simple extension of the funnel plot can help to facilitate an intuitive interpretation of the mathematics underlying a meta-analysis at a more fundamental level, by equating it to determining the centre of mass of a physical system. This analogy was recently exploited by MRC scientists to build a machine (named the 'Meta-Analyzer') to explain the concepts of weighing evidence and of biased evidence to a young audience at the Cambridge Science Festival, without recourse to precise definitions or statistical formulae. We have since produced a web-emulation of the machine to bring these ideas to an even wider audience of students and researchers - it can be found at <https://chjackson.shinyapps.io/MetaAnalyser/>.

Figure 1 shows the Meta-Analyzer web app populated with a fictional meta-analysis of $k=13$ studies. In it the standard funnel plot has been augmented so that the area representing point i is proportional to study i 's fixed effect weight $w_i = 1/s_i^2$, in order to promote its interpretation as a physical mass. It is proportional because, for study i , we show its weight in the Meta-Analyzer as a percentage of the total weight in the analysis, which equals $100 \times (w_i / \sum_{i=1}^k w_i)$.

Points are joined by horizontal cord to a pole that is joined itself to a vertical stand at a pivot point, p say. Study weights are imagined to exert a downward force due to gravity. Since the pole is perfectly

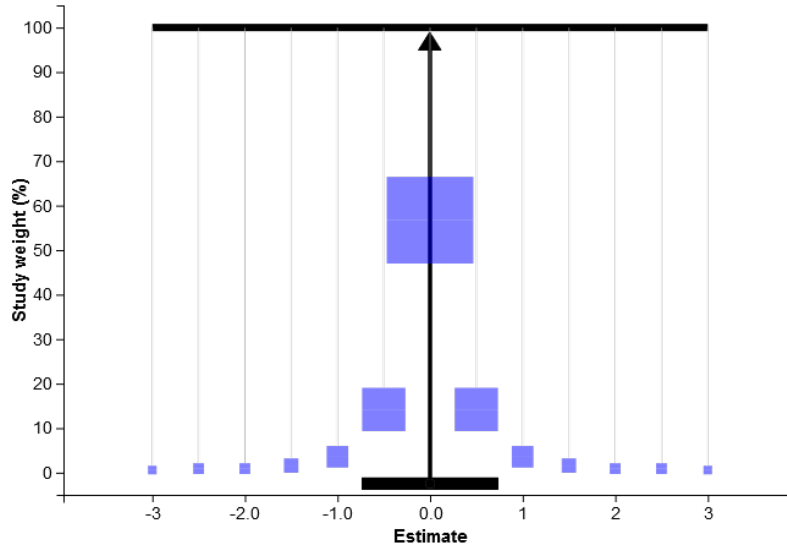


Fig. 1: The meta-analyzer supporting a fictional body of 13 study results. The centre of mass (overall estimate) is located at zero.

horizontal, it is intuitively understood that p satisfies the *physical law*:

$$\sum_{i=1}^k w_i(y_i - p) = 0,$$

and is therefore equal to the centre of mass (Beatty, 2005). The above formula can be viewed as a rudimentary estimating equation, a construction we continue to utilize throughout this paper. It is simple to verify that p is identical to the fixed effect estimate $\hat{\mu}$ in (2). The length of the stand base along the x-axis shows the 95% confidence interval for μ . Again, there is a nice physical analogy to draw. When there is a lot of uncertainty as to the overall effect estimate, so that $\text{Var}(\hat{\mu})$ is large, the stand length must be wide in order to properly account for its instability. Conversely, when the overall effect estimate is very precise, so that $\text{Var}(\hat{\mu})$ is small, the stand length need only be short. When populated with such idealised data as in Figure 1, the Meta-Analyzer may remind some of Galton’s bean board or quincunx, a tool that also exploits gravity to illustrate how statistical laws can emerge from a seemingly random physical process.

Although the Meta-Analyzer was intended to be a simple educational tool to demonstrate the basics of evidence synthesis, we have continued to find the physical system it describes useful more broadly, to easily explain some common secondary issues in meta-analysis and to help make new connections between statistical techniques for bias adjustment from the literature that seem, at first sight, unrelated. In Section 2 we show that it helps to transparently demonstrate the implications of moving from a fixed to a random effects model and to assess the influence of outlying studies. In Section 3 we discuss the issue of biased

evidence, how the Meta-Analyzer was used to explain this concept to a lay audience, and how small study bias is commonly addressed by medical statisticians using Egger regression (Egger et al, 1997). In Section 4, we review the method of Mendelian randomization – a technique for estimating the causal effect of a modifiable exposure on a health outcome using observational data, by circumventing the problem of confounding. We draw parallels between adjusting for bias in Mendelian randomization and small study bias in meta-analysis, and describe an extension to the standard model assumed by Egger regression first proposed in Bowden, Davey Smith and Burgess (2015) within the context of Mendelian randomization, that can (in theory) be applied to both fields. In Section 5 we show that, by viewing Egger regression from a causal inference perspective, a novel estimating equation interpretation of this method is found that can be intuitively visualized via the Meta-Analyzer. We illustrate this new interpretation on some real data examples and conclude with a discussion in Section 6.

2. Random effects models

A common issue in meta-analysis is accounting for between study heterogeneity. In a fixed effect meta-analysis all studies are assumed to provide an estimate of the same quantity, and the only difference between studies is in the precision of their respective estimates for this quantity. This is often thought to be an over-simplistic model, especially when flatly contradicted by Cochran’s Q-statistic, $Q = \sum_{i=1}^k \frac{1}{s_i^2} (y_i - \hat{\mu})^2$, being substantially larger than its expected value of $k-1$ under model (1), and a random effects model is preferred. Two distinct approaches for incorporating heterogeneity have emerged. The first, and most popular is via an additional additive random effect, as in model (3) below (e.g. DerSimonian and Laird, (1986), Higgins and Thompson, (2002)). The second is via the addition of a multiplicative scale factor, as in model (4) below (for example Thompson and Sharp (1999), Baker and Jackson (2013)).

$$y_i = \mu + s_i \epsilon_i + \delta_i, \quad \delta_i \sim N(0, \tau^2), \quad \epsilon_i \sim N(0, 1). \quad (3)$$

$$y_i = \mu + \phi^{\frac{1}{2}} s_i \epsilon_i, \quad \epsilon_i \sim N(0, 1). \quad (4)$$

We first note that, regardless of whether model (3) or (4) holds in practice, under the independence assumption $\perp\!\!\!\perp_{(\delta_i, s_i, \epsilon_i)}$, consistent estimation of the overall mean parameter μ is achieved by fitting the fixed effect model (1). In practice, estimation of μ can follow by simple application of formula (2) except now the weight given to study i , w_i , changes to $\frac{1}{s_i^2 + \tau^2}$ and $1/\phi s_i^2$ under models (3) and (4) respectively. The point estimate for μ obtained by fitting model (4) is identical to that obtained from fitting model (1). This is because the common term ϕ simply cancels from the numerator and denominator in (2) and only the variance of the estimate is altered. However, both the point estimate and variance for $\hat{\mu}$ change under the additive random effect model (3). Because its use is ubiquitous, especially in the field of medical

research (Baker and Jackson, 2013), we focus only on the additive model for the remainder of this section.

If, after fitting the fixed effect model (1), Cochran's Q statistic about $\hat{\mu}$ indicates additional heterogeneity ($Q > k - 1$), then it is common practice to estimate τ^2 using the procedure defined by DerSimonian and Laird (1986), to give $\hat{\tau}_{DL}^2$. This estimate conveniently provides a link between Q and a popular measure of heterogeneity, I^2 , (Higgins, 2002) as follows:

$$\frac{Q - (k - 1)}{Q} = \frac{\hat{\tau}_{DL}^2}{\hat{\tau}_{DL}^2 + s_{typ}^2} = I^2,$$

where s_{typ}^2 is referred to as the 'typical' within study variance.

2.1. Random effects models via the Meta-Analyzer

In order to provide a physical interpretation of the calculations underpinning a meta-analysis so that they can be implemented using the Meta-Analyser, we formulate a system of estimating equations (as a natural progression of the single estimating equation for fixed effect model (1)) to fit the random effects model to find μ and τ^2 as below:

$$\text{Weight equation: } w_i = 1/(s_i^2 + \tau^2) \quad (5)$$

$$\text{Mean equation: } \sum_{i=1}^k w_i(y_i - \mu) = 0 \quad (6)$$

$$\text{Heterogeneity equation: } \sum_{i=1}^k w_i(y_i - \mu)^2 - (k - 1) = 0. \quad (7)$$

Formula (7) is referred to as the generalized Q statistic (Bowden et al, 2011) and, when solved in conjunction with (5) and (6), it returns an estimate for μ and the Paule-Mandel (PM) estimate for τ^2 (Paule and Mandel, 1982), which we denote by $\hat{\tau}_{PM}^2$. As with $\hat{\tau}_{DL}^2$ the PM estimate is constrained to be positive, it is known to provide a more reliable estimate for the between study heterogeneity than $\hat{\tau}_{DL}^2$ (Veroniki et al, 2015). Random effects model (3) has been promoted by the Cochrane collaboration (Higgins and Green, 2011) and formally justified as a basis for inference beyond the current meta-analysis to future studies and populations (Higgins, Thompson and Spiegelhalter, 2009). It reduces to the fixed effects model when τ^2 is either fixed or estimated to be 0. This implies $\mu_i \equiv \mu$ for all i .

Application of the random effects model, with additional variance component τ^2 , leads to study results being both down-weighted and more similarly weighted. Furthermore, the original weight given to large studies is reduced to a greater extent than those of smaller studies. This issue is fairly subtle and

hard to comprehend, but the Meta-Analyzer provides a very simple visualization, by adjusting the mass of each weight. Figure 2 (left) shows the Meta-Analyzer populated with a data set of 8 randomized trial results, each assessing the use of magnesium to treat myocardial infarction. The effect measure, y_i , is the log-odds ratio of death between treatment and control groups for study i . These data were previously analyzed by Higgins and Spiegelhalter (2002b).

Between trial heterogeneity is present for these data ($\hat{\tau}_{DL}^2 = 0.095$, $I^2 = 27.6\%$). Under random effects model (3) w_i is reduced from $\frac{1}{s_i^2}$ to $1/(s_i^2 + \hat{\tau}^2)$. We represent the weight ‘loss’ induced by moving from a fixed to an additive random effects model by ‘drilling out’ a square of length x_i , in order to satisfy the pythagorean identity $x_i^2 + \frac{1}{s_i^2 + \hat{\tau}^2} = \frac{1}{s_i^2}$, as illustrated in Figure 3. The centre of mass defined by the holed-out weights and calculated by the Meta-Analyzer is automatically consistent with the random effects estimate for μ . In Figure 3 we plug in $\hat{\tau}_{DL}^2$ for $\hat{\tau}^2$ for illustration. Full results obtained via the estimating equation approach (and so using the PM estimate for τ^2) are shown in Table 1.

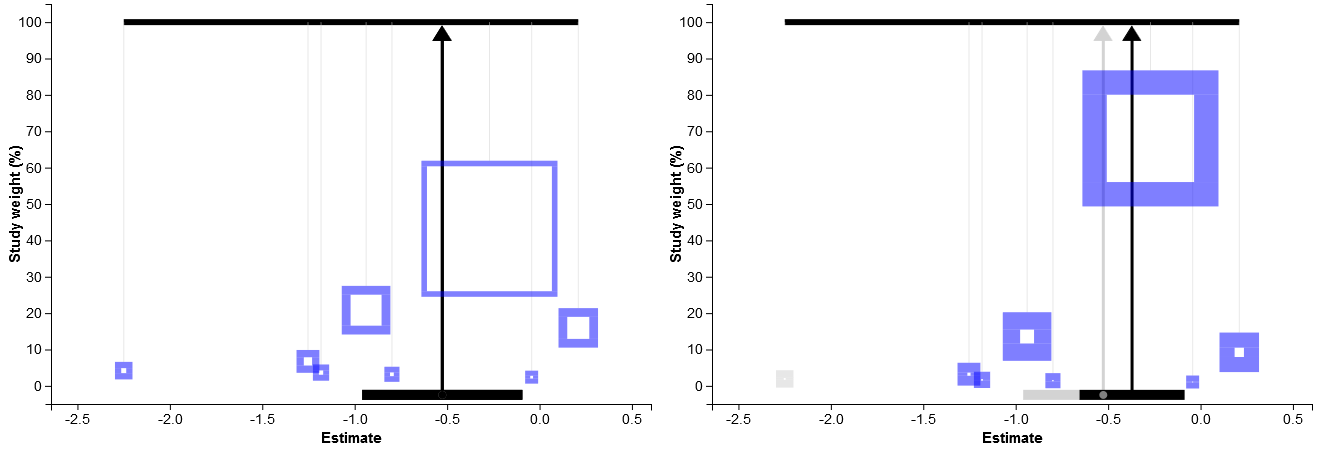


Fig. 2: *The Meta-Analyzer supporting the magnesium data under a random effects model for all trials (left); and with the Shechter trial removed (right).*

The fact that large studies lose more of their relative weight than small studies under an additive random effects model is immediately apparent from Figure 2 (left). We note briefly that no holing out is necessary when the multiplicative random effect model (4) is used. This is because the constant factor ϕ does not alter the weight given to study i as a proportion of the total weight in the analysis, whatever its value. However, when $\phi \neq 1$ then the variance of $\hat{\mu}$ will differ to that of the fixed effect model and so the stand length (confidence interval) will subsequently change.

2.2. Sensitivity to outliers

The amount of heterogeneity estimated in a meta-analysis can depend heavily on extreme, and often small, study results (Bowden et al, 2011). It is therefore useful in some circumstances to perform a sensi-

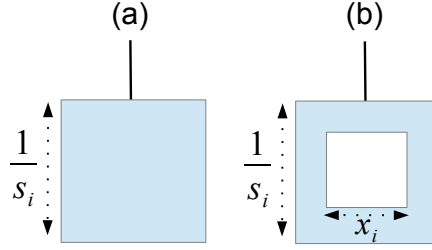


Fig. 3: (a) Weight given to study i in a fixed effect meta-analysis. (b) Weight given to study i in an additive random effects meta-analysis. A Square of length x_i is removed from weight i in order to satisfy the pythagorean formula.

tivity analysis, in which an outlying study result is excluded. Figure 2 (right) shows the Meta-Analyzer supporting the magnesium data under random effects model (3) excluding the Shechter study (shown in grey). The solid black support stand in Figure 2 (right) shows the overall estimate and corresponding 95% confidence interval in this case. The grey support stand shows the original point estimate and confidence interval. In this example, exclusion of the outlying study removes a large proportion of the between trial heterogeneity (updated $\hat{\tau}_{DL}^2 = 0.012$, $I^2 = 5.2\%$), making interpretation of the remaining trial data easier. Our web application facilitates easy transitions between various models like this as part of a sensitivity analysis. Users also see the Meta-Analyzer dynamically tip and re-balance in response to their latest analysis choice.

Model				
Parameter	Est	S.E	t value	p-value
All studies				
μ	-0.516	0.214	-2.408	0.047
τ_{PM}^2	0.084	-	-	-
τ_{DL}^2	0.095 ($I^2 = 27.6\%$)	-	-	-
Shechter study removed				
μ	-0.362	0.219	-1.653	0.149
τ_{PM}^2	0.008	-	-	-
τ_{DL}^2	0.012 ($I^2 = 5.1\%$)	-	-	-

Table 1: Meta-analysis of the Magnesium data under random effects model (3), with and without the Shechter trial.

3. Small study bias

3.1. The Aspirin data

Figure 4 (left) shows the Meta-Analyser enacted on 63 randomized controlled trials reported by Edwards et al. (1998) that each investigated the benefit of oral Aspirin for pain relief. Study estimates y_i represent

the log-odds ratios for the proportion of patients in each arm who had at least a 50% reduction in pain. Between trial heterogeneity was present for these data ($\tau_{DL}^2 = 0.04$, $I^2 = 10\%$) and Figure 4 (left) reflects the weight given to each study by the Meta-Analyzer under the random effects model (3) using τ_{DL}^2 as the heterogeneity parameter estimate. Despite the apparent between trial heterogeneity, the conclusion of the random effects meta-analysis is that oral Aspirin is an effective treatment, the combined log-odds ratio estimate is 1.26 in favour of Aspirin with a 95% confidence interval (1.1,1.41). Full results are shown in Table 2.

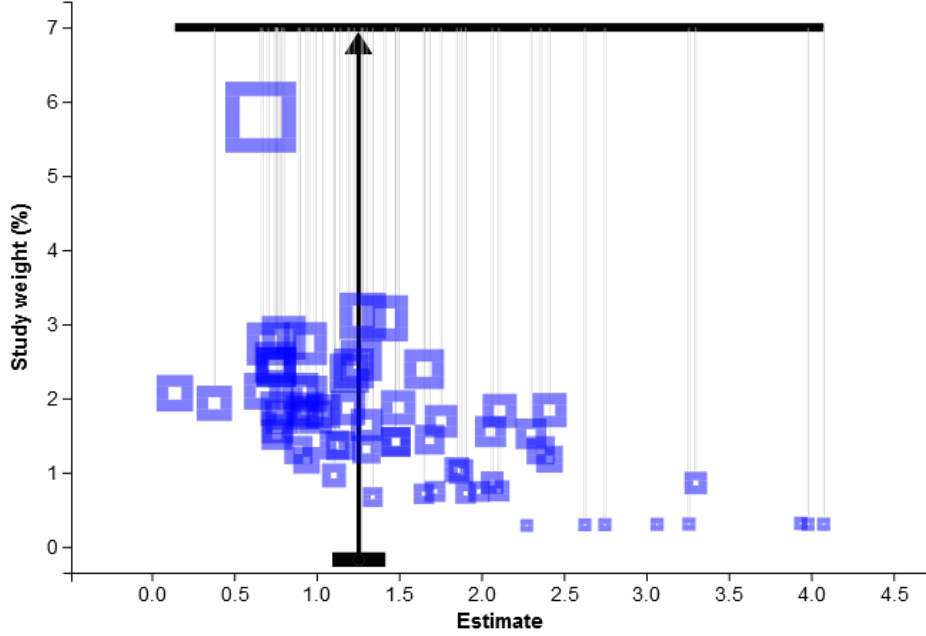


Fig. 4: *Meta-Analyzer supporting the Aspirin data under a random effects model.*

The hypothetical data shown in Figure 1 is perfectly symmetrical about its centre of mass, indicating that there is no correlation between effect size and precision across studies. However, there is a clear asymmetry present in the Aspirin data, smaller studies tend to show larger effect size estimates, whereas larger studies tend to report more modest results. For these data, $\text{Cor}(y_i, 1/s_i) = -0.7$, which suggests $\perp_{(\delta_i, s_i, \epsilon_i)}$ does not hold under model (3). The phenomenon of observing a negative correlation between study precision and effect size is often given the umbrella term ‘small study bias’ (Egger et al., 1997; Sterne et al., 2011; Rücker et al., 2011).

3.2. *The causes and consequences of small study bias*

Small study bias could actually be caused by real differences between small and large studies. Small trials may employ a more intensive intervention and therefore generate a greater effect on disease outcomes than larger trials (Egger et al, 1997; Bowater and Escrela, 2013). Asymmetry could also be a simple artefact of the data. For example, point estimates are not strictly independent of their estimated variances when

calculated from binary or count outcomes, (Harbord et. al, 2006, Peters et. al, 2010). However, its cause could be also be more sinister. Publication bias, or the file-drawer problem (Rosenthal, 1979) occurs when journals selectively publish study results that achieve a high level of statistical significance, and also induces asymmetry.

Much attention has been focused on methods to adjust for small study bias assuming that it is caused by selective dissemination of research findings, a practice which unfortunately is prevalent in biomedical research (Dwan et al, 2008). A common strategy is to propose an underlying selection model that lead to the generation of biased data. Many have followed Hedges (1984,1992), Copas (1999) and Copas and Shi (2000) in assuming the probability of observing a study result is some function of its precision. Whatever its true cause, when small study bias is present it can severely and adversely affect the conclusions reached. For example, it can lead one to detect apparent between trial heterogeneity when, in truth, none exists, and it can induce substantial bias into the overall estimate, $\hat{\mu}$, particularly under the random effects model (3), because it gives more relative weight to small studies than the fixed effect model (Henmi and Copas, 2010, Bowden et al, 2011). Indeed, for the Aspirin data, we see a slight reduction in the log-odds ratio estimate under the multiplicative random effects model (Table 2). In the case of the Aspirin data, it is reasonable to suspect that the effect of oral Aspirin on pain relief is substantially smaller than suggested by either random effects analysis.

3.3. *Small study bias explained to a lay audience via the Meta-Analyzer*

When attempting to explain the concept of biased evidence and of bias adjustment to the science festival audience, we opted for a simplified version of Trim and Fill (Duval and Tweedie, 2000). It aims to replace ‘missing studies’ in a meta-analysis by a process of reflection, until symmetry in the funnel plot is restored. We illustrated their idea within the context of a Sherlock Holmes’ style mystery (see Box 1 and Figure 9 in the Appendix which shows the Meta-Analyzer at its initial conceptual state and in situ at the Cambridge Science Festival ([\protect\vrulewidth0pthttp://www.sciencefestival.cam.ac.uk/](http://www.sciencefestival.cam.ac.uk/))).

3.4. *Egger regression*

Despite the intuition and appeal of its end result, the mathematics behind Trim and Fill are quite complicated. Perhaps due to its relative simplicity, the most popular approach to testing and adjusting for small study bias in medical research is Egger regression (Egger et al., 1997). This assumes the following linear fixed effect model in order to explain the correlation between y_i and $1/s_i$:

$$\frac{y_i}{s_i} = \beta_0 + \frac{\mu}{s_i} + \epsilon_i, \quad \epsilon_i \sim N(0, 1). \quad (8)$$

Model				
Parameter	Est	S.E	t value	p-value
Random effects model (3)				
μ	1.26	0.082	15.4	<10e-16
τ^2	0.04	-	-	-
Random effects model (4)				
μ	1.23	0.080	15.4	<10e-16
ϕ	1.11	-	-	-
Egger regression model (10)				
β_0	2.11	0.31	6.77	5.8e-09
μ	0.025	0.19	0.13	0.89
ϕ	0.64	-	-	-

Table 2: *Results for Meta-analyses of the Aspirin data.*

Testing for small study bias is then equivalent to testing $H_0: \beta_0 = 0$. If model (8) and $\perp\!\!\!\perp_{(s_i, \epsilon_i)}$ hold, then the overall effect estimate, $\hat{\mu}$, adjusted for possible small study bias (via $\hat{\beta}_0$) is a consistent estimate for the overall treatment effect, μ . Several authors have considered the addition of random effects into model (8), in order account for possible residual heterogeneity after adjustment for small study bias, see for example and Moreno et al, (2009), Peters et al, (2010) and Rücker et al, (2011). Their approaches have been straightforward generalisations of the additive and multiplicative random effects models (3) and (4) respectively, as below:

$$\frac{y_i}{s_i} = \beta_0 + \frac{\mu}{s_i} + \frac{\delta_i}{s_i} + \epsilon_i, \quad \delta_i \sim N(\mu, \tau^2) \quad \epsilon_i \sim N(0, 1). \quad (9)$$

$$\frac{y_i}{s_i} = \beta_0 + \frac{\mu}{s_i} + \phi^{\frac{1}{2}} \epsilon_i, \quad \epsilon_i \sim N(0, 1). \quad (10)$$

Again, regardless of whether model (9) or (10) holds in practice, under $\perp\!\!\!\perp_{(s_i, \delta_i, \epsilon_i)}$, consistent estimation of the overall mean parameter μ follows from fitting the standard fixed effect Egger model (8).

At first sight, model (9) seems the most natural extension to model (8). However, when a non-zero value for τ^2 is estimated under model (9), the resulting overall estimate for μ differs from, and can often exhibit more substantial bias than, the fixed effect estimate (Rücker et al. (2011)). By contrast, multiplicative model (10) is far more well behaved in this respect since its point estimate is identical to that obtained from fitting model (8). Nullifying the influence of variance components on the overall mean, a property enjoyed by model (10), is so attractive in the presence of small study bias, that approaches have been developed to artificially incorporate this feature into additive random effects models as well (Henmi and Copas, 2010).

For the reasons outlined above, we analyse the Aspirin data using the multiplicative Egger regression model only. This is straightforward, because model (10) is *automatically* fitted by the process of standard linear regression using the least squares criterion. For these data $\hat{\beta}_0 = 2.11$, with a p-value of approximately 1×10^{-8} , $\hat{\mu}$ is equal to 0.025, with a p-value of 0.89 and $\hat{\phi} = 0.64$. In summary, Egger regression detects a highly significant presence of small study bias and, after this has been removed, no evidence of a treatment effect whatsoever. A slightly concerning feature is the severe under-dispersion in the residual error after adjustment for small study bias, an issue we return to in Section 5.

4. Mendelian randomization and meta-analysis

Mendelian randomization (MR) (Davey Smith and Ibrahim, 2003) is a technique applied to observational data, that exploits information on genetic variants to infer whether a modifiable risk factor, X , has a causal effect on a health outcome, Y . Straightforward analysis methods with observational data are vulnerable to confounding bias, and estimates of association can only have a true causal interpretation if all confounders between X and Y (which we denote by U) have been adjusted for. An MR study can also be used in place of a randomized clinical trial, when such a trial can not be implemented for ethical, practical or financial reasons. The method requires that a genetic variant, G_i , exists that satisfies the ‘Instrumental Variable’ (IV) assumptions, see Bowden and Turkington (1990) for a thorough review. Specifically, G_i must be: (i) associated with X ; (ii) not associated U ; and (iii) not associated with Y given X and U . These assumptions are encoded by the solid arrows in Figure 5. The causal effect of X on Y is the parameter of interest and, following the notation in this paper, is denoted by μ . The Wald estimator

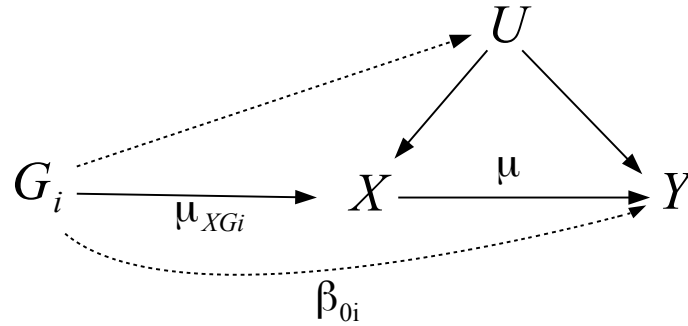


Fig. 5: Causal diagram representing the standard IV assumptions (solid lines), versus violations of the assumptions (dotted lines)

is the ratio of the gene i outcome association, $\hat{\mu}_{YG_i}$, and the gene i exposure association, $\hat{\mu}_{XG_i}$, giving $\hat{\mu}_i = \frac{\hat{\mu}_{YG_i}}{\hat{\mu}_{XG_i}}$. Under the IV assumptions, $\hat{\mu}_{YG_i}$ tends towards the product of the gene-exposure association and the causal effect, $\mu \hat{\mu}_{XG_i}$, as the sample size grows large so that the Wald estimate it is asymptotically unbiased.

Estimation of μ via a MR analysis incorporating multiple uncorrelated genetic variants G_1, \dots, G_k , can be viewed as equivalent to a fixed effect meta-analysis of the Wald ratio estimates $\hat{\mu}_1, \dots, \hat{\mu}_k$ (Johnson, 2014). So, formula (2) can be applied in the MR context to yield

$$\hat{\mu} = \frac{\sum_{i=1}^k \hat{\mu}_{XGi}^2 \hat{\mu}_i}{\sum_{i=1}^k \hat{\mu}_{XGi}^2} \quad (11)$$

Relating this back to the general meta-analysis context, we can equate the gene-exposure association, $\hat{\mu}_{XGi}$, to a within study precision, $1/s_i$, and the causal effect estimate, $\hat{\mu}_i$, with the effect estimate for study i , y_i . Just as for a fixed effects meta-analysis, formula (11) assumes that $\hat{\mu}_{XGi}$, although estimated, is known. It also assumes for simplicity that all variants have the same allele frequency, but this can easily be accounted for if false. Formula (11) has the advantage of being calculable directly from summary data estimates, not therefore requiring data at the individual level.

4.1. *Pleiotropy and small study bias*

In a short period, the number of genetic variants with summary data estimates available for a given MR study has gone from a handful to, in many cases, several hundred, thus fuelling an exponential growth in the technique's application and its statistical power. However, as the number of genetic variants utilized in a MR analysis increases, so too does the likelihood that some of the included variants are invalid instrumental variables, which could potentially bias the analysis. The dotted lines in Figure 5 highlight these possible violations, namely (ii) that a variant may be associated with a confounder of X and Y and (iii), that a variant may affect the outcome via a completely independent pathway than X . Violation of assumptions (ii) and (iii) can not be directly tested, and together the magnitude of their violation is referred to as the *pleiotropic* effect of a variant. Substituting $1/s_i$ for $\hat{\mu}_{XGi}$, y_i for $\hat{\mu}_i$ and denoting the total pleiotropic effect of variant i by β_{0i} , we can write the following more plausible model for the i 'th Wald ratio estimate in a MR analysis using the notation of this paper as follows:

$$\begin{aligned} \frac{y_i}{s_i} &= \beta_{0i} + \frac{\mu}{s_i} + \epsilon_i. \\ &= \beta_0 + \frac{\mu}{s_i} + \epsilon_i + \psi_i, \quad \epsilon_i \sim N(0, 1) \quad \psi_i \sim N(0, \sigma_{\beta_0}^2) \end{aligned} \quad (12)$$

Model (12), first conceptualized by Bowden, Davey Smith and Burgess (2015) but clarified here, allows the causal effect estimate of variant i to be composed of a common causal effect, μ , and an additional, individual pleiotropy term, $\beta_{0i} = \beta_0 + \psi_i$. For convenience we assume that ψ_i is normally distributed with zero mean and variance $\sigma_{\beta_0}^2$. The standard method of analysis for estimating the overall causal effect – formula (11) – assumes all genetic variants are valid IVs ($\beta_{0i} = 0$ for all i). Egger regression implemented via model (8) – which ignores the true heterogeneous nature of the bias by estimating only its mean value,

β_0 – can still consistently estimate the causal effect, μ , even when *all* genetic variants exhibit pleiotropy, as long as $\perp\!\!\!\perp_{(\psi_i, s_i, \epsilon_i)}$ holds for model (12). Under this assumption:

$$\hat{\mu} = \frac{\text{Cov}(y, \frac{1}{s})}{\text{Var}(1/s)} = \mu + \frac{\text{Cov}(\epsilon, \frac{1}{s})}{\text{Var}(1/s)} + \frac{\text{Cov}(\psi, \frac{1}{s})}{\text{Var}(1/s)}, \quad (13)$$

and as the number of variants k grows large, the covariance terms above tend to 0 and $\hat{\mu} \rightarrow \mu$. Bowden, Davey Smith and Burgess (2015) show that $\perp\!\!\!\perp_{(\psi_i, s_i, \epsilon_i)}$ is only plausible if IV assumptions (i) and (ii) hold. If assumption (ii) is violated (as indicated by the dotted line from G_i to the confounder U in Figure 5), ψ_i and $s_i = 1/\gamma_i$ will contain common terms via re-introduced confounding and will be positively correlated, thus Egger regression will no longer provide a consistent estimate for μ .

5. A causal interpretation of Egger regression

The application of Egger regression to a causal inference problem like Mendelian randomization is surprising, since this field has traditionally been dominated by instrumental variable methods that are conceptually very different. One such technique is the Structural Mean Model framework, which uses potential outcomes (Rubin, 2005) and places independence at the heart of its estimation strategy via the technique of G-estimation. For a general overview of this method see Vansteelandt and Joffe (2014) and, within the context of Mendelian randomization, Bowden and Vansteelandt, (2011). By making an analogy with Structural Mean Models, we now show that Egger regression can be understood as a means to de-bias a meta-analysis by restoring symmetry to the funnel plot, in a different but complimentary way to Trim and Fill. We start by assuming equation (10) as a ‘working’ mean model, even if in reality we believe that the heterogeneous bias model (12) is more realistic. We multiply each side of the working model by s_i and subtract $\beta_0 s_i$ to yield

$$y_i(\beta_0) = y_i - \beta_0 s_i = \mu + \phi^{\frac{1}{2}} \epsilon_i s_i.$$

The term $y_i(\beta_0)$ is a transformed version of the effect size estimates, can be viewed as a potential outcome, and is theoretically mean independent of s_i under $\perp\!\!\!\perp_{(s_i, \epsilon_i)}$. The intercept estimate obtained from fitting model (10), $\hat{\beta}_0$, can be viewed as defining a particular transform of the data $y_i(\hat{\beta}_0)$ that forces $y_i(\hat{\beta}_0)$ to be independent of s_i across all studies. Taking this further, model (10) is equivalent to solving the following system of estimating equations:

$$\text{Weight equation: } w_i = 1/(\phi s_i^2)$$

$$\text{Potential outcome transform: } y_i(\beta_0) = y_i - \beta_0 s_i$$

$$\text{Mean equation: } \sum_{i=1}^k w_i \{y_i(\beta_0) - \mu\} = 0 \quad (14)$$

$$\text{G-estimation equation: } \sum_{i=1}^k w_i \{y_i(\beta_0) - \mu\} (s_i - \bar{s}) = 0 \quad (15)$$

$$\text{Heterogeneity equation: } \sum_{i=1}^k w_i \{y_i(\beta_0) - \mu\}^2 - (k-2) = 0, \quad (16)$$

where \bar{s} is the arithmetic mean of the s_i terms. We now clarify the connection between the above system of estimating equations and estimation of β_0 , μ and ϕ using standard linear regression theory. Fitting model (8) to obtain estimates for β_0 and μ , is equivalent to solving equations (14) and (15) leaving ϕ unspecified (in the Appendix we provide some simple R code to verify this). We then formally define ϕ as a parameter and solve equation (16) to give

$$\hat{\phi} = \frac{\sum_{i=1}^k w_i \{y_i(\hat{\beta}_0) - \hat{\mu}\}^2}{k-2}. \quad (17)$$

We note the equivalence of the numerator of equation (17) to the Q' statistic defined in Rücker et al. (2011). The variances of $\hat{\beta}_0$ and $\hat{\mu}$ are given by

$$\text{Var}(\hat{\mu}) = \frac{\hat{\phi}}{\sum_{i=1}^k (1/s_i - 1/\bar{s})^2}, \quad \text{Var}(\hat{\beta}_0) = \text{Var}(\hat{\mu}) \bar{s}^2,$$

where \bar{s}^2 is the sample mean of the squared, within study variances.

Concerns over the use of Egger regression have been raised when analysing binary data because a study's outcome estimate will not truly be independent of its standard error, even when small study bias is not present. For this reason, Peters et al (2010) propose to replace s_i in the Egger regression equation with a measure of precision based on study size, $1/n_i$ say. This could easily be represented in our estimating equation with suitable modification. For example, the G-estimation formula above would then be used to find the β_0 that forces independence between $y_i(\beta_0) = y_i - \beta_0/n_i$ and $1/n_i$.

5.1. *Re-analysis of the Aspirin data*

Returning to the Aspirin data, Figure 6 (left) shows the final resting point of the Meta-Analyzer upon enacting Egger regression using the causal interpretation described above. Once the Egger regression

option is ticked, users observe a dynamical change in the Meta-Analyzer from its initial starting point of the standard random effects analysis in Figure 4. The x-axis position is now the transformed or potential outcome scale $y_i(\hat{\beta}_0) = y_i(2.11)$ for study i . The meta-analysis now exhibits a high amount of symmetry that can be immediately visualised by the user. This transformation is highlighted in Figure 6 (right), which plots potential outcomes $y_i(\hat{\beta}_0)$ on the horizontal versus $1/\hat{\phi}s_i$ on the vertical axis with the original data (y_i versus $1/s_i$) also shown for comparison. When small study bias is present, estimates from small studies are shifted by a relatively large amount in the horizontal dimension (in this case to the left), whereas those from large studies are shifted horizontally by a relatively small amount.

Because ϕ is estimated to be 0.64 in this instance, indicating under-dispersion after adjustment for small study bias, study precisions under the potential outcome transformation are increased in proportion to their size, so that the precisions of small studies are shifted vertically upwards by a small amount and those of large studies are shifted by a large amount. If over-dispersion had been observed after adjustment, we would have seen a shift vertically downwards instead. If we think that the heterogeneous bias model (12) is true and $\perp_{(\psi_i, s_i, \epsilon_i)}$ holds, then we would expect these data to be over-dispersed since, from equation (17), $E[\hat{\phi}] = 1 + \sigma_{\beta_0}^2$. Therefore, we can only meaningfully interpret $\hat{\phi} - 1$ as an estimate for $\sigma_{\beta_0}^2$ if $\hat{\phi} > 1$.

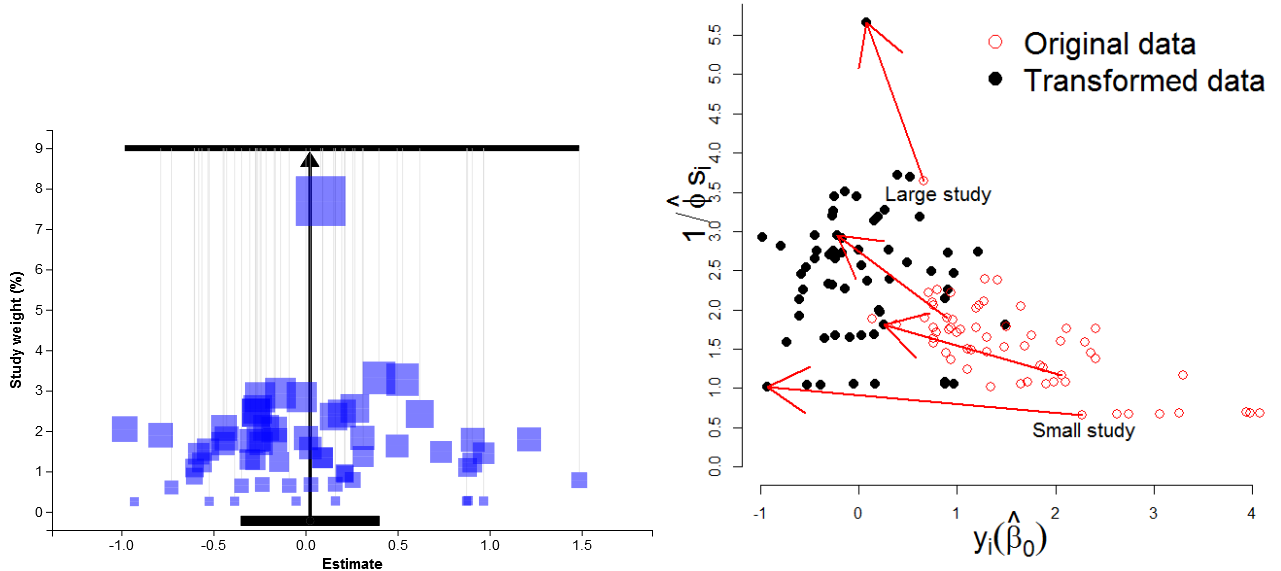


Fig. 6: Left: The Meta-Analyzer incorporating Egger regression enacted on the Aspirin data and shown under the potential outcome transform. Right: Funnel plot of the original Aspirin data (y_i vs $1/s_i$, hollow red dots) versus its transformed counterpart ($y_i(\hat{\beta}_0)$ vs $1/\hat{\phi}s_i$, solid black dots).

5.1.1. Can Egger regression correct for publication bias?

The heterogeneous bias model (12) seems well suited to the MR context. Furthermore, the work of Hedges and Copas *do* indeed imply that heterogeneous bias contaminates the outcome data, y_i , under publication bias. This is because the selective pressure study i experiences (and therefore the magnitude of bias in its estimate conditional on selection) is a function of its own, unique characteristics, such as s_i . So, whilst model (12) does therefore appear to be a realistic data generating model for mainstream meta-analyses affected by small study bias, if model (12) is true then $\perp\!\!\!\perp_{(\psi_i, s_i, \epsilon_i)}$ will be implausible because ψ_i will clearly then be correlated with s_i . Ironically, Egger regression is therefore unlikely to adequately adjust for small study bias when induced by outcome dependent sampling, such as in the clinical trial context (for which it was originally intended). Since under-dispersion is a phenomenon that can also easily be induced by publication bias (Copas and Shi, 2000) this provides even stronger evidence that the Egger regression analysis of the Aspirin data should be interpreted with extreme caution.

5.2. The causal effect of LDL cholesterol on heart disease risk

There is a long and extensive literature on the association between various lipid fractions and cardiovascular disease, but still far from universal agreement as to whether these associations have a causal basis. Using summary data estimates available on 154 variants from the CARDIOGRAM consortium (CARDIOGRAM, 2013), we perform a Mendelian randomization analysis to look for evidence that low density lipoprotein cholesterol (LDL-c) has a causal effect on coronary artery disease (CAD).

Since LDL-c levels are closely related and highly correlated with other lipid fractions, such as Triglycerides and high density lipoprotein, we selected only a subset of 57 variants out of the 154 that were most strongly associated LDL-c. The minimum p-value for the strength of association across all variants was $8.3\text{e-}7$. This strategy should reduce the possibility of violating causal assumption (ii), but does not rule it out completely. Furthermore, the selected variants might well exhibit pleiotropic effects through completely separate pathways, and therefore be in violation of assumption (iii). For this reason, we supplement the standard IVW analysis with Egger regression under model (12).

Figure 7 (left) and Table 3 shows the result of a standard additive random effects meta-analysis applied to the causal effect estimates across all 57 included variants. They quantify the causal effect in terms of a log-odds ratio of coronary artery disease for a 1 standard deviation increase in LDL-c levels. Significant heterogeneity is detected in the data ($I^2 = 70\%$, $\tau^2 = 0.11$), despite this, a strong positive causal log-odds ratio of 0.37 is estimated. However, there is reason to believe this analysis to be misleading, given that there exists a noticeable correlation between the magnitude of the causal effect estimate and its precision, indicative pleiotropy. Removal of the two relatively imprecise causal effect estimates

which are less than -1.5 reduces the heterogeneity considerably (adjusted $\tau^2=0.07$, results not shown).

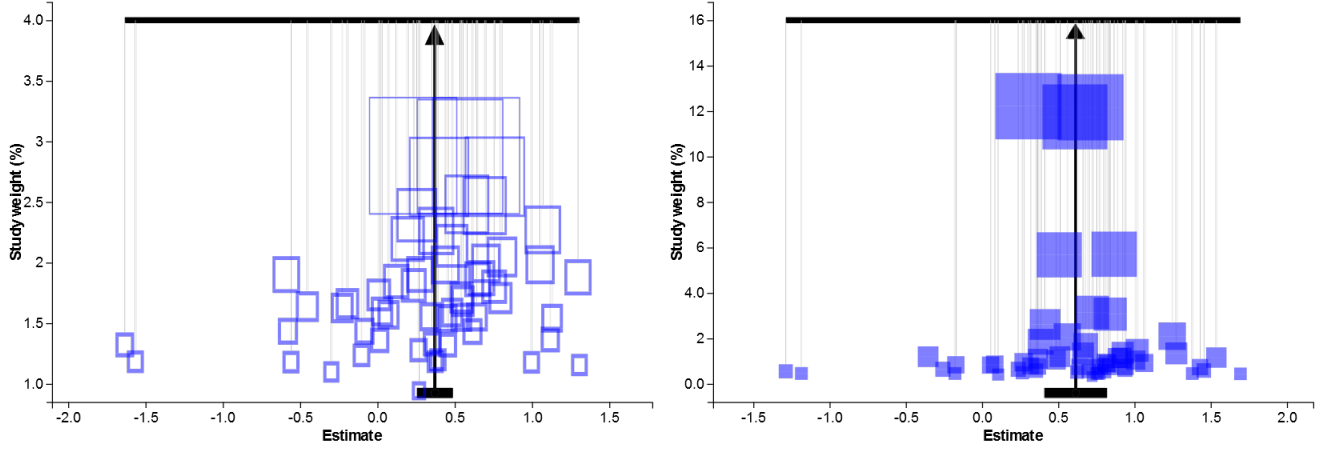


Fig. 7: *The Meta-Analyzer supporting an MR analysis of the LDL variants under an additive random effects model (left) and Egger regression (right).*

Model				
Parameter	Est	S.E	t value	p-value
Random effects model (3)				
μ	0.37	0.059	6.29	5.10e-08
τ^2	0.11	-	-	-
Random effects model (4)				
μ	0.45	0.053	8.51	1.13e-11
ϕ	3.33	-	-	-
Egger regression model (12)				
β_0	-0.0102	0.0046	-2.23	0.0298
μ	0.632	0.0975	6.481	2.66e-08
$\sigma_{\beta_0}^2$	2.11	-	-	-

Table 3: *MR-Analyses of the lipids data. Estimates for μ are log-odds ratios per standard deviation increase in LDL-c.*

Applying the multiplicative random effects model (4) to these data to remove the influence of the random effects variance, τ^2 , on the overall mean estimate, yields a slightly higher causal effect estimate of 0.45. Applying Egger regression model (12) – Figure 7 (right) and Table 3 – a significant negative effect of pleiotropy is detected, despite the pleiotropy variance being large ($\hat{\sigma}_{\beta_0}^2 = 2.11$). Consequently, the point estimate for μ is adjusted upward to 0.63. In conclusion, although significant evidence of pleiotropy exists across the included variants, there is still overwhelming evidence that LDL-c is causally related to CAD risk. If anything, the Egger analysis suggests that the true causal effect of LDL-c is slightly masked by pleiotropy acting in the opposing direction.

6. Discussion

In this paper we have shown that, by augmenting the funnel plot portraying a meta-analysis of study results with an additional pole, cord and pivot, we can give this abstract object – and the overall estimate that it implies – a clear physical interpretation. From its original conception and success as a science festival exhibit, we think the Meta-Analyzer has the continued potential to be a useful tool for educating and explaining the concept of weighing evidence and statistical reasoning to an even wider lay-audience, and hopefully to inspire the next generation of data scientists.

We have shown that the conceptual framework the Meta-Analyzer promotes is useful for both explaining the rationale, and interpreting the effect of, extended modelling choices in meta-analysis to a more technical, advanced audience. Finally, it has allowed us to make connections between methods to adjust for small study bias in meta-analysis, and for confounding bias in observational epidemiology.

It is worth commenting on the differing way in which heterogeneity is modelled in standard meta-analysis and causal inference. In the former, significant differences in effect estimates across studies (as measured by Cochran's Q statistic, say) is seen to provide clear evidence of underlying treatment effect heterogeneity and the subsequent adoption of a random effects model. However in causal inference it is strongly assumed that all instrumental variables identify a common causal effect. Any evidence of heterogeneity between causal effect estimates across different genetic instruments in an MR analysis is seen as evidence of pleiotropy (or invalid instruments more generally), and formal testing procedures such as the Sargan test (Sargan, 1951) exist to identify the invalid variants responsible so that they can be removed from the analysis. In our investigation of the magnesium data we did indeed remove an outlying study in order to demonstrate a dramatic reduction in the between trial heterogeneity. This is not encouraged within the general context of meta-analysis but perhaps, as in the causal inference field, it should be tolerated to a higher degree.

As the subject of Mendelian randomization moves forward - facilitated by the emergence of increasingly rich summary data sources, it is vital that the field makes use of existing methodology developed for meta-analysis, particularly in the area of bias modelling. However, by viewing the assumptions of established methods like Egger regression from a causal perspective, it is also possible that new insights can be gleaned and related back to the field of meta-analysis more generally. We hope that our paper can act as a catalyst to help promote and further this aim.

Acknowledgements

We Thank the MRC visual aids department for their help in building the Meta-Analyzer, the MRC Biostatistics Unit staff members who demonstrated its use at the Science fair and to Vikki O’neill in particular. We also thank Philip Haycock for help in accessing the CARDIoGRAM data, as well as Dan Jackson and George Davey Smith for their helpful comments on an early draft of this work. Dr Jack Bowden is supported by an MRC methodology research fellowship (grant number MR/L012286/1).

References

- Baker, R. and Jackson, D. (2013). Meta-analysis inside and outside particle physics: two traditions that should converge? *Research Synthesis Methods*, **4**, 109–124.
- Beatty, M. (2005). *Principles of Engineering Mechanics, Volume 2: Dynamics: The Analysis of Motion*. Springer Science & Business Media.
- Bowater, R. and Escarela, G. (2013). Heterogeneity and study size in random-effects meta-analysis. *Journal of Applied Statistics*, **40**, 2–16
- Burgess, S., Thompson, S. (2013). Bayesian methods for meta-analysis of causal relationships estimated using genetic instrumental variables *Statistics in medicine*, **29**, 1298–1311
- Bowden, J., and Vansteelandt, S. (2011) Mendelian randomization analysis of case-control data using structural mean models *Statistics in Medicine*, **30**, 678–694.
- Bowden, J., Tierney, J. Copas, A., and Burdett, S. (2011). Quantifying, displaying and accounting for heterogeneity in the meta-analysis of RCTs using standard and generalised Q statistics. *Medical Research Methodology*, **11**, 41.
- Bowden, J. Davey Smith, G., and Burgess, S. (2015). Mendelian randomization with invalid instruments: effect estimation and bias detection through Egger regression *International Journal of Epidemiology*, **44**, 512–25.
- Bowden, R. J., and Turkington, D. A. (1990) Instrumental variables. Vol. 8. Cambridge University Press.
- The CARDIoGRAMplusC4D Consortium. (2013). Large-scale association analysis identifies new risk loci for coronary artery disease. *Nature Genetics*, **45**, 25–33.
- Copas, J. (1999). What works? Selectivity models in meta-analysis. *JRSSA*, **162**, 95–109.
- Copas, J. and Shi, J. (2000). Meta-analysis, funnel plots and sensitivity analysis. *Biostatistics*, **1**, 247–262

- Davey Smith G., and Ebrahim S. (2003). Mendelian randomization: can genetic epidemiology contribute to understanding environmental determinants of disease? *International Journal of Epidemiology*, **32**, 1–22.
- DerSimonian, R. and Laird, N. (1986). Meta-analysis in clinical trials. *Controlled Clinical Trials*, **7**, 177–188.
- Didelez, V., Meng, S., Sheehan, N. A. (2010). Assumptions of IV methods for observational epidemiology. *Statistical Science* **25**, 22–40.
- Duval, S. and Tweedie, R (2000). Trim and fill: A simple funnel-plot-based method of testing and adjusting for publication bias in meta-analysis. *Biometrics*, **56**, 455–463.
- Dwan, K., Altman, D. G., Arnaiz, J. A, Bloom, J., Chan, A. W., Cronin, E., Decullier, E., Easterbrook, P. J., Von Elm, E., Gamble, C., Ghera, D., Ioannidis, J. P. A., Simes, J., Williamson, P. R., (2008). Systematic review of the empirical evidence of study publication bias and outcome reporting bias. *PLoS ONE*, **3**, e3081.
- Edwards, J. E. Oldman, A., Smith, L., Collins, S. L., Carol, D., Wiffen, P. J., McQuay, H.J., and Moore, R.A. (1998) Single dose oral aspirin for acute pain. *Cochrane Database of Systematic Reviews*, **4**.
- Egger, M., Davey Smith, G., Schneider, M., and Minder, C. (1997) Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal*, **315**, 629–634.
- Harbord, R. M., Egger, M., and Sterne, J. (2006). A modified test for small-study effects in meta-analyses of controlled trials with binary endpoints. *Statistics in Medicine*, **25**, 3443–3457
- Higgins, J. P. T. and Thompson, S. (2002). Quantifying heterogeneity in a meta-analyses. *Statistics in Medicine*, **21**, 1539–1558.
- Hedges, L. V. (1992). Modeling publication selection effects in meta-analysis. *Statistical Science*, **7**: 246–255.
- Henmi, M. and Copas, J. (2010). Confidence intervals for random effects meta-analysis and robustness to publication bias. *Statistics in Medicine*, **29**, 2969–2983.
- Higgins, J. P. T. and Spiegelhalter, D. J. (2002b). Being sceptical about meta-analyses: a Bayesian perspective on magnesium trials in myocardial infarction. *International Journal of Epidemiology*, **31**, 96104, 2002.
- Higgins, J. P. T., Thompson, S. G., and Spiegelhalter, D.J. (2009). A re-evaluation of random-effects meta-analysis. *J. Royal Statistical Soc. Series A*, **172**, 137–159.

- Higgins, J. P. T., and Green, S. (2011). *Cochrane Handbook for Systematic Reviews of Interventions* Version 5.1.0. The Cochrane Collaboration.
- Johnson, T. (2014). Efficient calculation for multi-SNP genetic risk scores. Technical Report, The Comprehensive R Archive Network 2013. Available at <http://cran.r-project.org/web/packages/gtx/vignettes/ashg2012.pdf> [last accessed 2014/11/19].
- Moreno, S. G., Sutton, A. J., Ades, A. E., Stanley, T. D., Abrams, K. R., Peters, J. L., Cooper, N. J. (2009). Assessment of regression-based methods to adjust for publication bias through a comprehensive simulation study *BMC Medical Research Methodology*, **9**, 2.
- Peters, J. L., Sutton, A. J., Jones, D. R., Abrams, K. R., Rushton, L., Moreno, S. G. (2010). Assessing publication bias in meta-analyses in the presence of between study heterogeneity. *JRSSA*, **173**, 575–591.
- Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, **86** , 638–641.
- Paule, R. and Mandel, J. (1982). Concensus values and weighting factors. *J Res Natl Bur Stand*, **87**, 377–85.
- Rubin, D.B. (2005). Causal inference using potential outcomes *Journal of the American Statistical Association*, **100**, 322–331 .
- Rücker, G., Schwarzer, G., Carpenter, J., Binder, H., and Schumacher, M. (2011). Treatment-effect estimates adjusted for small study effects via a limit meta-analysis. *Biostatistics*, **12**, 122–142.
- Sterne, J. and Egger, M (2001). Funnel plots for detecting bias in meta-analysis: Guidelines on choice of axis. *Journal of Clinical Epidemiology*, **54**, 1046–1055.
- Sterne, J., Sutton, A., Ioannidis, J., et al. (2011). Recommendations for examining and interpreting funnel plot asymmetry in meta-analyses of randomised controlled trials. *British Medical Journal*, **343**, d4002
- Thompson, S. and Sharp, S. (1999). Explaining heterogeneity in meta-analysis: a comparison of methods. *Statistics in Medicine*, **18**, 2693–2708.
- Vansteelandt, S. and Joffe, M. (2014). Structural nested models and g-estimation: The partially realized promise. *Statistical Science*, **29**, 707–731.
- Veroniki, A., Jackson, D., Viechtbauer, W., Bender, R., Bowden, J., Knapp, G., Kuss, O., Higgins, JPT, Langan, D., Salanti, G. (2015) Methods to estimate the between-study variance and its uncertainty in meta-analysis *Research Synthesis Methods*, To appear

The Meta-Analyzer from theory to practice

Box 1: Sherlock Holmes and the case of the missing evidence

‘It is 1858. A new machine has been invented to exploit the ethereal force of gravity to weigh evidence: The Meta-Analyzer. Its inventor, Stata Lovelace says “*The Meta-Analyser can make decisions affecting Britain and her empire in the name of gold-standard fairness and morality*”. The night before its maiden calculation (into the affect of sanitation on patient mortality: is it beneficial?), Professor Moriarty has stolen some evidence! Call Scotland Yard! And Sherlock Holmes!’

‘Inspector Lestrade inspects the lop-sided Meta-Analyzer (Figure 8 (left)) and mops his brow’

Inspector Lestrade: ‘It’s simple ’olmes, just move the pivot point’ (Figure 8 (middle)), ‘Case closed.’

Sherlock Holmes: ‘Its a solution, but an inelegant one, Lestrade. Replace the missing studies to return the balance and yield the unbiased truth’ (Figure 8 (right))

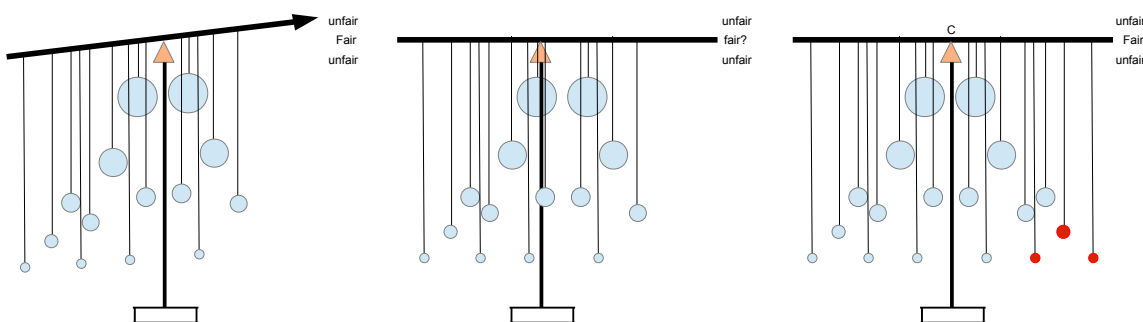


Fig. 8: Left: *Unbalanced Meta-Analyzer*. Middle: *Rebalanced Meta-Analyzer (Inspector Lestrade approach)*. Right: *Rebalanced Meta-Analyzer (Sherlock Holmes approach)*

R code for Section 5

```

## given Aspirin data vectors y,s
## G-estimation routine
G_est = function(a){
  w      = 1/s^2
  Beta0  = a[1]
  yBeta0 = y - Beta0*s
  MU     = sum(w*yBeta0)/sum(w)
  L      = (sum(w*(yBeta0-MU)*(s - mean(s))))^2
}
Beta0hat = optimize(G_est,c(-5,5))$min
yBeta0hat = y - Beta0hat*s
w      = 1/s^2
mu     = sum(w*yBeta0hat)/sum(w)
>Beta0hat
[1] 2.112803
> mu
[1] 0.02519816
> ## standard Egger regression
> summary(lm(y~s,weights=1/s^2))$coef[,1]
(Intercept)      s
0.02519816  2.11280317

```



Fig. 9: *The Meta-Analyzer in action at the Cambridge Science Festival*